Modeling Trolling in Social Media Conversations

Luis Gerardo Mojica
Human Language Technology Research Institute
The University of Texas at Dallas
Richardson, TX 75083-0688
mojica@utdallas.edu

arXiv:1612.05310v2 [cs.CL] 28 Dec 2016

Abstract

Social media websites, electronic newspapers and Internet forums allow visitors to leave comments for others to read and interact. This exchange is not free from participants with malicious intentions, who *troll* others by positing messages that are intended to be provocative, offensive, or menacing. With the goal of facilitating the computational modeling of trolling, we propose a trolling categorization that is novel in the sense that it allows comment-based analysis from both the trolls' and the responders' perspectives, characterizing these two perspectives using four aspects, namely, the troll's intention and his intention disclosure, as well as the responder's interpretation of the troll's intention and her response strategy. Using this categorization, we annotate and release a dataset containing excerpts of Reddit conversations involving suspected trolls and their interactions with other users. Finally, we identify the difficult-to-classify cases in our corpus and suggest potential solutions for them.

1 Introduction

In contrast to traditional content distribution channels like television, radio and newspapers, Internet opened the door for direct interaction between the content creator and its audience. Young people are now gaining more frequent access to online, networked media. Although most of the time, their Internet use is harmless, there are some risks associated with these online activities, such as the use of social networking sites (e.g., Twitter, Facebook, Reddit). The anonymity and freedom provided by social networks makes them vulnerable to threatening situations on the Web, such as *trolling*.

Trolling is "the activity of posting messages via a communications network that are intended to be provocative, offensive or menacing" (Bishop, 2013). People who post such comments are known as *trolls*. According to Hardaker (2010), a troll's "real intention(s) is/are to cause disruption and/or trigger or exacerbate conflict for the purpose of their own amusement". Worse still, the troll's comments may have a negative psychological impact on his target/victim and possibly others who participated in the same conversation. It is therefore imperative to identify such comments and perhaps even terminate the conversation before it evolves into something psychological disruptive for the participants. Monitoring conversations is a labor-intensive task: it can potentially place a severe burden on the moderators, and it may not be an effective solution when traffic is heavy. This calls for the need to develop automatic methods for identifying malicious comments, which we will refer to as *trolling attempts* in this paper.

In fact, there have recently been some attempts to automatically identify comments containing *cyberbullying* (e.g., Van Hee et al. (2015)), which corresponds to the most severe cases of trolling (Bishop, 2013). However, we believe that it is important not only to identify trolling attempts, but also comments that could have a negative psychological impact on their recipients. As an example, consider the situation where a commenter posts a comment with the goal of amusing others. However, it is conceivable that not everybody would be aware of these playful intentions, and these people may disagree or dislike the mocking comments and take them as inappropriate, prompting a negative reaction or psychological impact on themselves.

In light of this discussion, we believe that there is a need to identify not only the trolling attempts, but also comments that could have a negative psychological impact on its receipts. To this end, we seek to

achieve the following goals in this paper. First, we propose a comprehensive categorization of trolling that allows us to model not only the troll's intention given his trolling attempt, but also the recipients' *perception* of the troll's intention and subsequently their reaction to the trolling attempt. This categorization gives rise to very interesting problems in pragmatics that involve the computational modeling of intentions, perceived intentions, and reactions to perceived intentions. Second, we create a new annotated resource for computational modeling of trolling. Each instance in this resource corresponds to a *suspected* trolling attempt taken from a Reddit conversation, it's surrounding context, and its immediate responses and will be manually coded with information such as the troll's intention and the recipients' reactions using our proposed categorization of trolling. Finally, we identify the instances that are difficult to classify with the help of a classifier trained with features taken from the state of the art, and subsequently present an analysis of these instances.

To our knowledge, our annotated resource is the first one of its sort that allows computational modeling on both the troll's side and the recipients' side. By making it publicly available, we hope to stimulate further research on this task. We believe that it will be valuable to any NLP researcher who is interested in the computational modeling of trolling.

2 Related Work

In this section, we discuss related work in the areas of trolling, bullying, abusive language detection and politeness, as they intersect in their scope and at least partially address the problem presented in this work.

In the realm of psychology, Bishop (2013) and Bishop (2014) elaborate a deep description of a troll's personality, motivations, effects on the community that trolls interfere in and the criminal and psychological aspects of trolls. Their main focus are flaming (trolls), and hostile and aggressive interactions between users (Osullivan and Flanagin, 2003).

On the computational side, Mihaylov et al. (2015) address the problem of identifying manipulation trolls in news community forums. Not only do they focus solely on troll identification, but the major difference with this work is that all their predictions are based on non-linguistic information such as number of votes, dates, number of comments and so on. In a networks related framework, Kumar et al. (2014) and Guha et al. (2004) present a methodology to identify malicious individuals in a network based solely on the network's properties rather than on the textual content of comments. Cambria et al. (2010) propose a method that involves NLP components, but fail to provide an evaluation of their system.

There is extensive work on detecting offensive and abusive language in social media (Nobata et al., 2016) and (Xiang et al., 2012). There are two clear differences between their work and ours. One is that trolling is concerned about not only abusive language but also a much larger range of language styles and addresses the intentions and interpretations of the commenters, which goes beyond the linguistic dimension. The other is that we are additionally interested in the reactions to trolling attempts, real or perceived, because we argued that this is a phenomenon that occurs in pairs through the interaction of at least two individuals, which is different from abusive language detection. Also, Xu et al. (2012a), Xu et al. (2012b) and Xu et al. (2013) address bullying traces. Bullying traces are self-reported events of individuals describing being part of bullying events, but we believe that the real impact of computational trolling research is not on analyzing retrospective incidents, but on analyzing real-time conversations. Chen et al. (2012) use lexical and semantic features to determine sentence offensiveness levels to identify cyberbullying, offensive or abusive comments on Youtube. On Youtube as well, Dinakar et al. (2012) identified sensitive topics for cyberbullying. Dadvar et al. (2014) used expert systems to classify between bullying and no bullying in posts. Van Hee et al. (2015) predict fine-grained categories for cyberbullying, distinguishing between insults and threats and identified user roles in the exchanges. Finally, Hardaker (2010) argues that trolling cannot be studied using established politeness research categories.

3 Trolling Categorization

In this section, we describe our proposal of a comprehensive trolling categorization. While there have been attempts in the realm of psychology to provide a working definition of trolling (e.g., Hardaker (2010), Bishop (2014)), their focus is mostly on modeling the troll's behavior. For instance, Bishop (2014) constructed a "trolling magnitude" scale focused on the severity of abuse and misuse of internet mediated communications. Bishop (2013) also categorized trolls based on psychological characteristics focused on pathologies and possible criminal behaviors. In contrast, our trolling categorization seeks to model not only the troll's behavior but also the impact on the recipients, as described below.

Since one of our goals is to *identify* trolling events, our datasets will be composed of *suspected* trolling attempts (i.e., comments that are suspected to be trolling attempts). In other words, some of these suspected trolling attempts will be real trolling attempts, and some of them won't. So, if a suspected trolling attempt is in fact not a trolling attempt, then its author will *not* be a troll.

To cover both the troll and the recipients, we define a (suspected trolling attempt, responses) pair as the basic unit that we consider for the study of trolling, where "responses" are all the direct responses to the suspected trolling attempt. We characterize a (suspected trolling attempt, responses) pair using four aspects. Two aspects describe the trolling attempt: (1) **Intention (I)** (what is its author's purpose?), and (2) **Intention Disclosure (D)** (is its author trying to deceive its readers by hiding his real (i.e., malicious) intentions?). The remaining two aspects are defined on each of the (direct) responses to the trolling attempt: (1) **Intention Interpretation (R)** (what is the responder's perception of the troll's intention?), and (2) the **Response strategy (B)** (what is the responder's reaction?). Two points deserve mention. First, R can be different from I due to misunderstanding and the fact that the troll may be trying to hide his intention. Second, B is influenced by R, and the responder's comment can itself be a trolling attempt. We believe that these four aspects constitute interesting, under-studied pragmatics tasks for NLP researchers.

The possible values of each aspect are described in Table 1. As noted before, since these are *suspected* trolling attempts, if an attempt turns out not to be a trolling attempt, its author will not be a troll.

For a given (suspected trolling attempt, responses) pair, not all of the 189 ($= 3 \times 3 \times 3 \times 7$) combinations of values of the four aspects are possible. There are logical constraints that limit plausible combinations: a) *Trolling* or *Playing Intentions* (I) must have *Hidden* or *Exposed* Intention Disclosure (D), b) *Normal* intentions (I) can only have *None* Intention disclosure (D) and c) *Trolling* or Playing interpretation (R) cannot have *Normal* response strategy (B).

3.1 Conversation Excerpts

To enable the reader to better understand this categorization, we present two example excerpts taken from the original (Reddit) conversations. The first comment on each excerpt, generated by author C0, is given as a minimal piece of context. The second comment, written by the author C1 in italics, is the suspected trolling attempt. The rest of the comments comprise all direct responses to the suspected trolling comment.

Example 1.

C0: Yeah, cause that's what usually happens. Also, quit following me around, I don't want a boyfriend.
  C1: *I wasn't aware you were the same person.... I've replied to a number of stupid people recently, my bad*
    C0: Trollname trollpost brotroll

In this example, C1 is teasing of C0, expecting to provoke or irritate irritate, and he is clearly *disclosing* her trolling intentions. In C0's response, we see that he clearly believe that C1 is trolling, since is directly calling him a "brotroll" and his *response strategy* is frustrate the trolling attempt by denouncing C1 troll's intentions "trollpost" and true identity "brotroll".

Example 2.

C0: Please post a video of your dog doing this. The way I'm imagining this is adorable.
  C1: I hope the dog gets run over by a truck on the way out of the childrens playground.
    C2: If you're going to troll, can you at least try to be a bit more
    C0: Haha I hope the cancer kills you. convincing?

In this example, we observe that C0's first comment is making a polite request (*Please*). In return, C1 answer is a mean spirited comment whose *intention* is to disrupt and possible hurtful C0. Also, C1's

Class	Description	Size %
Intention (I)		
Trolling	The comment is malicious in nature, aims to disrupt, annoy, offend, harm or spread purposely false information	53.6% (537)
Mock Trolling or Playing	The comment is playful, joking, teasing or mocking others without the malicious intentions as in the Trolling class	8.9% (89)
No Trolling	A simple comment without malicious or playful intentions.	37.7% (375)
Intention Disclosure (D)		
Exposed	A troll, clearly exposing its malicious or playful intentions	34.7% (347)
Hidden	A troll hiding its real malicious or playful intentions	11.5% (115)
None	The comment's author is not a troll, therefore there are no hidden nor exposed malicious or playful intentions	53.8% (539)
Intentions Interpretation (R)		
Trolling	The responder believes that the suspected troll is being malicious, annoying, offensive, harmful or attempts to spread false information	59.7% (785)
Mock Trolling or Playing	The responder believes that the suspected troll is being playful, joking, teasing or mocking without the malicious intentions	5.3% (70)
No Trolling	The responder believes that the suspected comment has no malicious intentions nor is playful, it is a simple comment.	35.0% (461)
Response Strategy (B)		
Engage	Fall in the perceived provocation, showing an emotional response, upset or annoyed	26.9% (354)
Praise	Acknowledge the perceived malicious or playful intentions and positively recognize the troll's ingenuity or ability	3.0% (39)
Troll	Acknowledge the perceived malicious and counter attack with a trolling attempt	24.0% (316)
Follow	Acknowledge the perceived malicious or playful intention and play along with the troll, further trolling	3.0% (39)
Frustrate	Acknowledges the perceived malicious or playful intentions and attempt to criticize or minimize them	13.0% (171)
Neutralize	Acknowledges the perceived malicious or playful intentions and give no importance to them	9.5% (125)
Normal	There is no perception or interpretation of a trolling attempt and the response is a standard comment	20.7% (272)

Table 1: Classes for trolling aspects: Intention, Intention Disclosure, Intention Interpretation and Response Strategy. *Size* refers to the percentage per class, in parenthesis is the total number of instances in the dataset.

comment is not subtle at all, so his intention is clearly *disclosed*. As for C2, she is clearly acknowledging C1's trolling intention and her *response strategy* is a criticism which we categorize as *frustrate*. Now, in C0's second comment, we observe that his *interpretation* is clear, he believes that C1 is *trolling* and the negative effect is so tangible, that his response strategy is to *troll* back or *counter-troll* by replying with a comparable mean comment.

4 Corpus and Annotation

Reddit[1] is popular website that allows registered users (without identity verification) to participate in fora grouped by topic or interest. Participation consists of posting stories that can be seen by other users, voting stories and comments, and comments in the story's comment section, in the form of a forum. The forums are arranged in the form of a tree, allowing nested conversations, where the replies to a comment are its direct responses. We collected all comments in the stories' conversation in *Reddit* that were posted in August 2015. Since it is infeasible to manually annotate all of the comments, we process this dataset with the goal of extracting threads that involve suspected trolling attempts and the direct responses to them. To do so, we used Lucene[2] to create an inverted index from the comments and queried it for comments containing the word "**troll**" with an edit distance of 1 in order to include close variations of this word, hypothesizing that such comments would be reasonable candidates of real trolling attempts. We did observe, however, that sometimes people use the word *troll* to point out that another user is trolling. Other times, people use the term to express their frustration about a particular user, but there is no trolling attempt. Yet other times people simply discuss trolling and trolls without actually observing one. Nonetheless, we found that this search produced a dataset in which 44.3% of

[1]https://www.reddit.com/

[2]https://lucene.apache.org/

the comments are real trolling attempts. Moreover, it is possible for commenters to believe that they are witnessing a trolling attempt and respond accordingly even where there is none due to misunderstanding. Therefore, the inclusion of comments that do not involve trolling would allow us to learn what triggers a user's interpretation of trolling when it is not present and what kind of response strategies are used.

For each retrieved comment, we reconstructed the original conversation tree it appears in, from the original post (i.e., the root) to the leaves, so that its parent and children can be recovered[3]. We consider a comment in our dataset a suspected trolling attempt if at least one of its immediate children contains the word *troll*. For annotation purposes, we created snippets of conversations exactly like the ones shown in *Example 1* and *Example 2*, each of which consists of the parent of the suspected trolling attempt, the suspected trolling attempt, and all of the direct responses to the suspected trolling attempt.

We had two human annotators who were trained on snippets (i.e., (suspected trolling attempt, responses) pairs) taken from 200 conversations and were allowed to discuss their findings. After this training stage, we asked them to independently label the four aspects for each snippet. We recognize that this limited amount of information is not always sufficient to recover the four aspects we are interested in, so we give the annotators the option to discard instances for which they couldn't determine the labels confidently. The final annotated dataset consists of 1000 conversations composed of 6833 sentences and 88047 tokens. The distribution over the classes per trolling aspect is shown in the table 2 in the column "Size".

Due to the subjective nature of the task we did not expect perfect agreement. However, on the 100 doubly-annotated snippets, we obtained substantial inter-annotator agreement according to Cohen's kappa statistic (Cohen, 1968) for each of the four aspects: Intention: 0.788, Intention Disclosure: 0.780, Interpretation: 0.797 and Response 0.776. In the end, the annotators discussed their discrepancies and managed to resolve all of them.

5 Trolling Attempt Prediction

In this section, we make predictions on the four aspects of our task, with the primary goal of identifying the errors our classifier makes (i.e., the hard-to-classify instances) and hence the directions for future work, and the secondary goal of estimating the state of the art on this new task using only shallow (i.e., lexical and wordlist-based) features.

5.1 Feature Sets

For prediction we define two sets of features: (1) a basic feature set taken from Van Hee's (2015) paper on cyberbullying prediction, and (2) an extended feature set that we designed using primarily information extracted from wordlists and dictionaries.

5.1.1 Basic Feature Set

N-gram features. We encode each lemmatized and unlemmatized unigram and bigram collected from the training comments as a binary feature. In a similar manner, we include the unigram and bigram along with their POS tag as in (Xu et al., 2012a). To extract these features we used Stanford CoreNLP (Manning et al., 2014).

Sentiment Polarity. The overall comment's emotion could be useful to identify the response and intention in a trolling attempt. So, we apply the Vader Sentiment Polarity Analyzer (Hutto and Gilbert, 2014) and include four features, one per each measurement given by the analyzer: positive, neutral, negative and a composite metric, each as a real number value.

5.1.2 Extended Feature Set

Emoticons. Reddit's comments make extensive use of emoticons. We argue that some emoticons are specifically used in trolling attempts to express a variety of emotions, which we hypothesize would be useful to identify a comment's intention, interpretation and response. For that reason, we use the emoticon dictionary developed Hogenboom et al. (2015). We create a binary feature whose value is one if at least one of these emoticons is found in the comment.

[3]We removed the comments whose text had been deleted

Harmful Vocabulary. In their research on bullying, Nitta et al. (2013) identified a small set of words that are highly offensive. We create a binary feature whose value is one if the comment contains at least one of these words.

Emotions Synsets. As in Xu et al. (2012b), we extracted all lemmas associated with each WordNet (Miller, 1995) *synset* involving seven emotions (anger, embarrassment, empathy, fear, pride, relief and sadness) as well as the synonyms of these emotion words extracted from the English Merriam-Webster (2004) dictionary. We create a binary feature whose value is one if any of these synsets or synonyms appears in the comment.

Swearing Vocabulary. We manually collected 1061 swear words and short phrases from the internet, blogs, forums and smaller repositories [4]. The informal nature of this dictionary resembles the type of language used by flaming trolls and agitated responses, so we encode a binary feature whose value is one when at least one such swear word is found in the comment.

Swearing Vocabulary in Username. An interesting feature that is suggestive of the intention of a comment is the author's username. We found that abusive and annoying commenters contained cursing words in their usernames. So, we create a binary feature whose value is one if a swear word from the swearing vocabulary is found in their usernames.

Framenet. We apply the SEMAFOR parser (Das et al., 2014) to each sentence in every comment, and construct three different types of binary features: every frame name that is present in the sentence, the frame name and the target word associated with it, and the argument name along with the token or lexical unit in the sentence associated with it. We believe that some frames are especially interesting from the trolling perspective. We hypothesize that these features are useful for identifying trolling attempts in which semantic and not just syntactic information is required.

Politeness cues. Danescu-Niculescu-Mizil et al. (2013) identified cues that signal polite and impolite interactions among groups of people collaborating online. Based on our observations of trolling examples, it is clear that flaming, hostile and aggressive interactions between users (Osullivan and Flanagin, 2003) and engaged or emotional responses would use impolite cues. In contrast, neutralizing and frustrating responses to the troll avoid falling in confrontation and their vocabulary tends to be more polite. So we create a binary feature whose value is one if at least one cue appears in the comment.

GloVe Embeddings. All the aforementioned features constitute a high dimensional bag of words (BOW). Word embeddings were created to overcome certain problems with the BOW representation, like sparsity, and weight in correlations of semantically similar words. For this reason, and following Nobata et al. (2016), we create a distributed representation of the comments by averaging the word vector of each lowercase token in the comment found in the Twitter corpus pre-trained GloVe vectors (Pennington et al., 2014). The resulting comment vector representation is a 200 dimensional array that is concatenated with the existing BOW.

5.2 Results

Using the features described in the previous subsection, we train four independent classifiers using logistic regression[5], one per each of the four prediction tasks. All the results are obtained using 5-fold cross-validation experiments. In each fold experiment, we use three folds for training, one fold for development, and one fold for testing. All learning parameters are set to their default values except for the regularization parameter, which we tuned on the development set. In Table 2 the leftmost results column reports F1 score based on majority class prediction. The next section (*Single Feature Group*) reports F1 scores obtained by using one feature group at a time. The goal of the later set of experiments is to gain insights about feature predictive effectiveness. The right side section (*All features*) shows the system performance measured using recall, precision, and F-1 as shown when all features described in section 5.1 are used.

The majority class prediction experiment is simplest baseline to which we can can compare the rest of the experiments. In order to illustrate the prediction power of each feature group independent from

[4] http://www.noswearing.com/

[5] We use the *scikit-learn*(Pedregosa et al., 2011) implementation

all others, we perform the "Single Feature Group", experiments. As we can observe in Table 2 there are groups of features that independently are not better than the majority baseline, for example, the emoticons, politeness cues and polarity are not better *disclosure* predictors than the majority base. Also, we observe that only n-grams and GloVe features are the only group of features that contribute to more than a class type for the different tasks. Now, the "All Features" experiment shows how the interaction between feature sets perform than any of the other features groups in isolation. The accuracy metric for each trolling task is meant to provide an overall performance for all the classes within a particular task, and allow comparison between different experiments. In particular, we observe that GloVe vectors are the most powerful feature set, accuracy-wise, even better than the experiments with all features for all tasks except interpretation.

		Single Feature Group										All Features			
Aspect/Class	mjr	emt	hrm	syn	swr	usr	frm	cue	pol	ngr	glv	R	P	F1	Size
I: Intention															
No trolling	69.7	69.4	69.2	68.3	67.3	66.2	69.6	69.6	69.6	64.2	71.9	76.0	60.8	67.5	53
Trolling	0.0	1.0	2.9	4.2	32.2	7.4	0.0	0.0	0.0	41.2	27.4	41.8	49.0	45.1	38
Playing	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	9
Accuracy	53.5	53.4	53.4	52.5	54.6	51.1	53.5	53.5	53.5	52.8	57.9	-	-	56.4	-
D: hidden															
Hidden	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12
Exposed	0.0	0.0	0.5	0.0	32.5	5.7	0.0	0.0	0.0	51.2	31.8	49.8	53.0	49.8	35
None	69.9	69.3	69.5	69.6	70.3	69.2	69.7	69.7	69.7	66.5	72.0	76.3	61.5	76.3	53
Accuracy	53.9	53.4	53.6	53.5	57.5	53.1	53.9	53.9	53.9	56.8	58.9	-	-	57.9	-
R: Interpretation															
No trolling	0.0	4.0	0.5	1.0	4.7	0.0	0.0	0.0	0.0	50.8	38.7	63.0	56.3	63.0	34
Playing	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	6
Trolling	74.2	73.6	74.1	73.8	74.2	74.2	74.2	74.2	74.2	74.3	76.3	74.3	72.8	74.3	60
Accuracy	59.2	58.4	59.2	58.5	59.2	59.2	59.2	59.2	59.2	64.4	64.9	-	-	66.3	-
R: Interpretation															
Frustrate	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	30.0	19.4	30.3	28.0	28.0	12
Troll	0.0	2.7	0.9	4.5	28.4	0.0	0.0	0.0	0.0	33.6	27.3	36.3	34.5	34.5	24
Follow	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	18.8	6.1	66.8	14.5	14.5	3
Praise	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	5.8	0.0	25.0	3.8	3.8	3
Neutralize	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	37.0	26.8	38.3	43.8	43.8	9
Normal	0.0	2.8	0.0	0.0	1.0	0.0	0.0	0.0	0.0	41.2	35.4	37.5	42.3	42.3	20
Engage	43.8	44.5	43.9	43.2	42.7	43.7	43.7	43.7	43.7	40.5	49.9	42.0	45.5	45.5	29
Accuracy	36.0	28.4	27.9	27.9	30.1	28.1	28.1	28.1	28.1	36.0	38.0	-	-	37.5	-
All Tasks Combined															
Total Accuracy	52.5	48.4	48.5	48.1	50.4	47.9	48.7	48.7	48.7	52.5	54.9	-	-	54.5	-

Table 2: *Experiments Results*. Below the "mjr" header, we report F1 scores the the majority class prediction we report F1 scores for the four aspects of trolling: Intention, Intentions Disclosure, Interpretation, and Response strategy. Also, below the "Single Feature Group" header, we report F1 scores as before, when the feature group indicated in the column headers is the only feature group used for classifier. The column headers abbreviations stand for: Emoticons, Harmful Vocabulary, Emotion Synsets, Swearing Vocabulary, Swearing Vocabulary in Usernames, Framenet, Politeness cues, n-grams (actual n-grams and n-grams appended with their corresponding part of speech tag) and Glove embeddings in that order. Below the "All Features" header we report Recall, Precision and F1 score, respectively, when all features are use for prediction. All experiments are performed using a logistic regression classifier per task. The last column reports the class distribution in percentage per task. The last row of each trolling aspect reports accuracy (the percentage of instances correctly classified). The last row in the table reports total accuracy, the percentage of correctly classified instances considering all aspects.

The overall Total Accuracy score reported in table 2 using the entire feature set is 549. This result is what makes this dataset interesting: there is still lots of room for research on this task. Again, the primary goal of this experiment is to help identify the difficult-to-classify instances for analysis in the next section.

6 Error Analysis

In order to provide directions for future work, we analyze the errors made by the classifier trained on the extended features on the four prediction tasks.

Errors on *Intention* (I) prediction: The **lack of background** is a major problem when identifying trolling comments. For example, "your comments fit well in Stormfront" seems inoffensive on the surface. However, people who know that Stormfront is a white supremacist website will realize that the author of this comment had an annoying or malicious intention. But our system had no knowledge about it and simply predicted it as non-trolling. These kind of errors reduces recall on the prediction of trolling comments. A solution would be to include additional knowledge from anthologies along with a sentiment or polarity. One could modify NELL (Mitchell et al., 2015) to broaden the understanding of entities in the comments.

Non-cursing aggressions and insults This is a challenging problem, since the majority of abusive and insulting comments rely on profanity and swearing. The problem arises with subtler aggressions and insults that are equally or even more annoying, such as "Troll? How cute." and "settle down drama queen". The classifier has a more difficult task of determining that these are indeed aggressions or insults. This error also decreases the recall of trolling intention. A solution would be to exploit all the comments made by the suspected troll in the entire conversation in order to increase the chances of finding curse words or other cues that lead the classifier to correctly classify the comment as trolling.

Another source of error is the presence of **controversial topic words** such as "black","feminism", "killing", "racism", "brown", etc. that are commonly used by trolls. The classifier seems too confident to classify a comment as trolling in the presence of these words, but in many cases they do not. In order to ameliorate this problem, one could create ad-hoc word embeddings by training glove or other type of distributed representation on a large corpus for the specific social media platform in consideration. From these vectors one could expect a better representation of controversial topics and their interactions with other words so they might help to reduce these errors.

Errors on *Disclosure* (D) prediction: A major source of error that affects disclosure is the **shallow meaning representation** obtained from the BOW model even when augmented with the distributional features given by the glove vectors. For example, the suspected troll's comment "how to deal with refugees? How about a bullet to the head" is clearly mean-spirited and is an example of disclosed trolling. However, to reach that conclusion the reader need to infer the meaning of "bullet to the head" and that this action is desirable for a vulnerable group like migrants or refugees. This problem produces low recall for the disclosed prediction task. A solution for this problem may be the use of deeper semantics, where we represent the comments and sentences in their logical form and infer from them the intended meaning.

Errors on *Interpretation* (R) prediction: it is a common practice from many users to **directly ask** the suspected troll if he/she is trolling or not. There are several variations of this question, such as "Are you a troll?" and "not sure if trolling or not". While the presence of a question like these seems to give us a hint of the responder's interpretation, we cannot be sure of his interpretation without also considering the context. One way to improve interpretation is to exploit the response strategy, but the response strategy in our model is predicted independently of interpretation. So one solution could be similar to the one proposed above for the disclosure task problem: jointly learning classifiers that predict both variables simultaneously. Another possibility is to use the temporal sequence of response comments and make use of older response interpretation as input features for later comments. This could be useful since commenters seem to influence each other as they read through the conversation.

Errors on *Response Strategy* (B) prediction: In some cases there is a **blurry line between "Frustrate" and "Neutralize"**. The key distinction between them is that there exists some *criticism* in the *Frustrate* responses towards the suspected troll's comment, while "Neutralizing" comments acknowledge that the suspected troll has trolling intentions, but gives no importance to them. For example, response comments such as "oh, you are a troll" and "you are just a lame troll" are examples of this subtle difference. The first is a case of "neutralize" while the second is indeed criticizing the suspected troll's comment and therefore a "frustrate" response strategy. This kind of error affects both precision and recall for these two classes. A possible solution could be to train a specialized classifier to disambiguate between "frustrate"

and "neutralize" only.

Another challenging problem is the **distinction between the classes "Troll" and "Engage"**. This is true when the direct responder is intensely flared up with the suspected comment to the point that his own comment becomes a trolling attempt. A useful indicator for distinguishing these cases are the presence of insults, and to detect them we look for swear words, but as we noted before, there is no guarantee that swear words are used for insulting. This kind of error affects the precision and recall for the "troll" and "engage" classes. A solution to this problem may be the inclusion of longer parts of the conversation. It is typical in a troll-engaged comment scheme to observe longer than usual exchanges between two users, and the comments evolve in very agitated remarks. One may then use this information to disambiguate between the two classes.

7 Conclusion and Future Work

We presented a new view on the computational modeling of trolling in Internet fora where we proposed a comprehensive categorization of trolling attempts that for the first time considers trolling from not only the troll's perspective but also the responders' perspectives. This categorization gives rise to four interesting pragmatics tasks that involve modeling intensions, perceived intensions, and reactions. Perhaps most importantly, we create an annotated dataset that we believe is the first of its sort. We intend to make publicly available with the hope of stimulating research on trolling.

References

[Bishop2013] Jonathan Bishop. 2013. The effect of de-individuation of the internet troller on criminal procedure implementation: An interview with a hater. *International Journal of Cyber Criminology*, 7(1):28.

[Bishop2014] Jonathan Bishop. 2014. Representations of trolls in mass media communication: a review of media-texts and moral panics relating to internet trolling. *International Journal of Web Based Communities*, 10(1):7–24.

[Cambria et al.2010] Erik Cambria, Praphul Chandra, Avinash Sharma, and Amir Hussain. 2010. Do not feel the trolls. *ISWC, Shanghai*.

[Chen et al.2012] Ying Chen, Yilu Zhou, Sencun Zhu, and Heng Xu. 2012. Detecting offensive language in social media to protect adolescent online safety. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*, pages 71–80. IEEE.

[Cohen1968] Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4):213.

[Dadvar et al.2014] Maral Dadvar, Dolf Trieschnigg, and Franciska de Jong. 2014. Experts and machines against bullies: a hybrid approach to detect cyberbullies. In *Advances in Artificial Intelligence*, pages 275–281. Springer.

[Danescu-Niculescu-Mizil et al.2013] Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. A computational approach to politeness with application to social factors. *arXiv preprint arXiv:1306.6078*.

[Das et al.2014] Dipanjan Das, Desai Chen, André FT Martins, Nathan Schneider, and Noah A Smith. 2014. Frame-semantic parsing. *Computational Linguistics*, 40(1):9–56.

[Dinakar et al.2012] Karthik Dinakar, Birago Jones, Catherine Havasi, Henry Lieberman, and Rosalind Picard. 2012. Common sense reasoning for detection, prevention, and mitigation of cyberbullying. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 2(3):18.

[Guha et al.2004] Ramanthan Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. 2004. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM.

[Hardaker2010] Claire Hardaker. 2010. Trolling in asynchronous computer-mediated communication: From user discussions to academic definitions. *Journal of Politeness Research*, 6(2):215–242.

[Hogenboom et al.2015] Alexander Hogenboom, Daniella Bal, Flavius Frasincar, Malissa Bal, Franciska De Jong, and Uzay Kaymak. 2015. Exploiting emoticons in polarity classification of text. *J. Web Eng.*, 14(1&2):22–40.

[Hutto and Gilbert2014] Clayton J Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International AAAI Conference on Weblogs and Social Media*.

[Kumar et al.2014] Sudhakar Kumar, Francesca Spezzano, and VS Subrahmanian. 2014. Accurately detecting trolls in slashdot zoo via decluttering. In *Advances in Social Networks Analysis and Mining (ASONAM), 2014 IEEE/ACM International Conference on*, pages 188–195. IEEE.

[Manning et al.2014] Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.

[Merriam-Webster2004] Merriam-Webster. 2004. *Merriam-Webster's collegiate dictionary*. Merriam-Webster.

[Mihaylov et al.2015] Todor Mihaylov, Georgi D Georgiev, AD Ontotext, and Preslav Nakov. 2015. Finding opinion manipulation trolls in news community forums. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning, CoNLL*, volume 15, pages 310–314.

[Miller1995] George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.

[Mitchell et al.2015] T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. 2015. Never-ending learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*.

[Nitta et al.2013] Taisei Nitta, Fumito Masui, Michal Ptaszynski, Yasutomo Kimura, Rafal Rzepka, and Kenji Araki. 2013. Detecting cyberbullying entries on informal school websites based on category relevance maximization. In *IJCNLP*, pages 579–586.

[Nobata et al.2016] Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th International Conference on World Wide Web*, pages 145–153. International World Wide Web Conferences Steering Committee.

[Osullivan and Flanagin2003] Patrick B Osullivan and Andrew J Flanagin. 2003. Reconceptualizing flamingand other problematic messages. *New Media & Society*, 5(1):69–94.

[Pedregosa et al.2011] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

[Pennington et al.2014] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

[Van Hee et al.2015] Cynthia Van Hee, Els Lefever, Ben Verhoeven, Julie Mennes, Bart Desmet, Guy De Pauw, Walter Daelemans, and Véronique Hoste. 2015. Detection and fine-grained classification of cyberbullying events. In *International Conference Recent Advances in Natural Language Processing (RANLP)*, pages 672–680.

[Xiang et al.2012] Guang Xiang, Bin Fan, Ling Wang, Jason Hong, and Carolyn Rose. 2012. Detecting offensive tweets via topical feature discovery over a large scale twitter corpus. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 1980–1984. ACM.

[Xu et al.2012a] Jun-Ming Xu, Kwang-Sung Jun, Xiaojin Zhu, and Amy Bellmore. 2012a. Learning from bullying traces in social media. In *Proceedings of the 2012 conference of the North American chapter of the association for computational linguistics: Human language technologies*, pages 656–666. Association for Computational Linguistics.

[Xu et al.2012b] Jun-Ming Xu, Xiaojin Zhu, and Amy Bellmore. 2012b. Fast learning for sentiment analysis on bullying. In *Proceedings of the First International Workshop on Issues of Sentiment Discovery and Opinion Mining*, page 10. ACM.

[Xu et al.2013] Jun-Ming Xu, Benjamin Burchfiel, Xiaojin Zhu, and Amy Bellmore. 2013. An examination of regret in bullying tweets. In *HLT-NAACL*, pages 697–702.